THE SCIENCE OF INFORMATION PROTECTION

Daniel F. Warren

Naval Postgraduate School

Abstract: The presentation of Information Protection material can be improved in two important ways. First, if the material is arranged in a systematic/scientific fashion it can show how all the various pieces fit together and it can also demonstrate completeness by showing that all threats are addressed. Second, if each protection technique is preceded by a clear description of the threat that it addresses learning is significantly enhanced because the protection technique is motivated. This paper presents an information threat model that 1) arranges the material in a scientific/systematic fashion and 2) facilitates a threat-first presentation of Information Protection techniques.

Key words: Information Protection, Access Control

1. INTRODUCTION

The threat model presented here is based on the following seven threat classes. These threat classes are mutually disjoint and together constitute all threats.

1. The Non-Human threat
2. The Human Error threat
3. The Authorized Person threat
4. The Unauthorized Insider threat
5. The Outsider Attacking Enterprise Data on the LAN threat
6. The Outsider Attacking Enterprise Data Outside the LAN threat
7. The Malicious Software threat

1.1 The following sections describe:

- why most Information Protection treatments fail to be scientific/systematic,
- why most Information Protection treatments fail to motivate the material,
- the nature of the threat model,
- how the model can be used to arrange Information Protection material in a scientific/systematic fashion,
- how the model facilitates a threat-first presentation of Information Protection techniques and
- why the classes of the threat model constitute all threats.

2. WHY CURRENT INFORMATION PROTECTION TREATMENTS FAIL TO BE SCIENTIFIC/SYSTEMATIC

A recent ad for a new information security book announces that the book covers the following major themes: (sub topics are also given for each major theme)

Cryptography, **Access Control**, **Protocols** and **Software**

A scientific minded reader is bound to ponder:

- Do these topics cover all facets of Information Protection?
- What are the relationships between these different themes?
- Is there a significance to the given order?

In general, Information Security or Information Protection books present topics in a somewhat ad hoc manner that does not lend itself to answering these questions. This, in turn, lessens their effectiveness as teaching tools. At a minimum every reader would like to know if an Information Protection book covers all aspects of information protection.

The ability to answer this question and the others posed above requires a systematic approach, which is a direct consequence of following the scientific method. The Merriam-Webster on-line dictionary[1] defines "scientific method" as:

> "**:** principles and procedures for the systematic pursuit of knowledge involving the recognition and formulation of a problem, ... "

The lack of a scientific approach is noteworthy, since, we are, after all, computer scientists. The periodic table of elements is a good example of a systematic arrangement of information that provides answers to these questions. It does list all currently known elements, it does show the relationships between different classes of elements (Alkali metals, Alkali earth metals, Transition metals, Rare earth metals, Noble gases and Halogens) and there is a significance to the order of elements in the table.

The result of this systematic arrangement of elements is a framework that facilitates an understanding of their properties.

The threat model presented here systematically arranges the various topics of Information Protection. Thus, it facilitates learning in the same way that the periodic table of elements facilitates learning.

3. WHY CURRENT TREATMENTS FAIL TO MOTIVATE THE MATERIAL

Information Protection and First Aid are similar in the sense that they are both problem/solution disciplines. If a First Aid course is organized around solutions the approach is less than optimal and fails to motivate the material. Such a presentation might start with aspirin and describe what aspirins do and what aspirins don't do. It might then present bandages and describe what they do and what they don't do. And so forth.

This material is better presented if it is organized around the problems. For acid indigestion do this and this and this. For bleeding do this and this and this. And so forth. When presented this way the student's interest is first peaked by the problem. "Hmm, how is acid indigestion addressed?" Then when the solution is presented it has context and meaning. The student is waiting to hear the solution and is likely to better comprehend it because of the context.

Currently most Information Protection books are arranged around solutions, like Cryptography, Access Control, Protocols, etc. And often the context is ignored. When execution domains are presented how often is the threat that they address also presented? If the student does not know what the problem is they are not going to be tuned into the solution. From a pedagogical standpoint a better approach is to arrange Information Protection topics around the problems and then present the protection techniques that address them.

Since the threat model introduced here arranges Information Protection material around the problems it naturally motivates an engaging presentation of protection techniques.

4. THE THREAT MODEL

The threat model presented here is based on a collection of seven threat classes that are mutually disjoint and together constitute all threats. These classes are based upon a typical enterprise computing situation, like the one shown below. This situation includes two Local Area Networks (LANs) that are physically far apart so they communicate via a Wide Area Network (WAN), like the Internet. The goal of the enterprise Information Protection policy is to protect enterprise information from unauthorized observation, modification and denial of availability and to protect the authenticity of source information for enterprise messages.

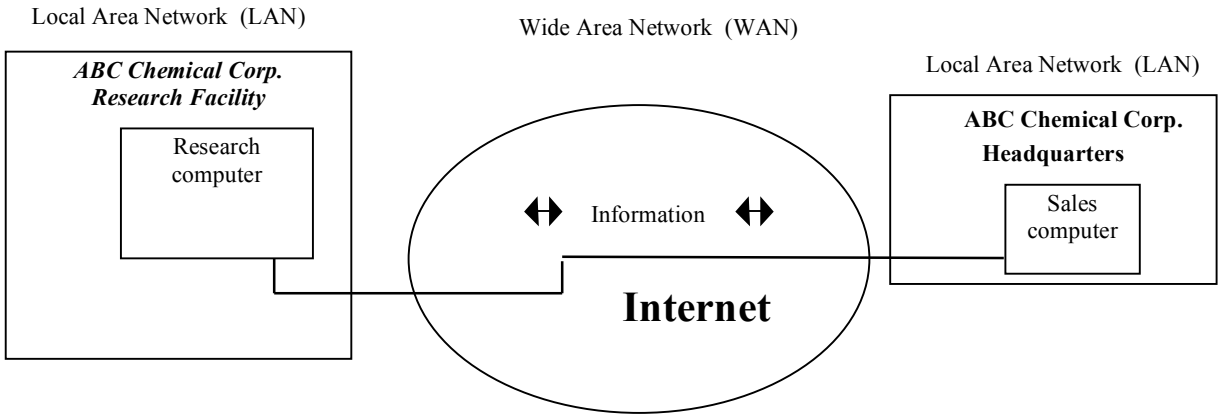

Figure 1. Typical Enterprise Computing Environment

Persons that are authorized to access enterprise information (such as employees) are called "insiders." Persons that are not authorized to access enterprise information are called "outsiders." The LAN systems and network infrastructure are assumed to be protected by typical LAN physical and personnel security measures.

This means that information residing on the LAN systems or on the LAN wired networks does not need the same type of protection that is required for information that is residing on systems and networks outside the LAN or in a wireless portion of the LAN.

The model refers to information residing on the LAN systems or on the LAN wired networks as "***information on the LAN***" and information that is residing on systems and networks outside the LAN or in a wireless portion of the LAN as "***information that is outside the LAN***."

The seven threat classes of the model are:

(1) Non-Human threat
(2) Human Error threat
(3) Authorized Person threat
(4) Unauthorized Insider threat
(5) Outsider Attacking Enterprise Data on the LAN threat
(6) Outsider Attacking Enterprise Data Outside the LAN threat
(7) Malicious Software threat (orchestrated by participants of classes (4) or (5))

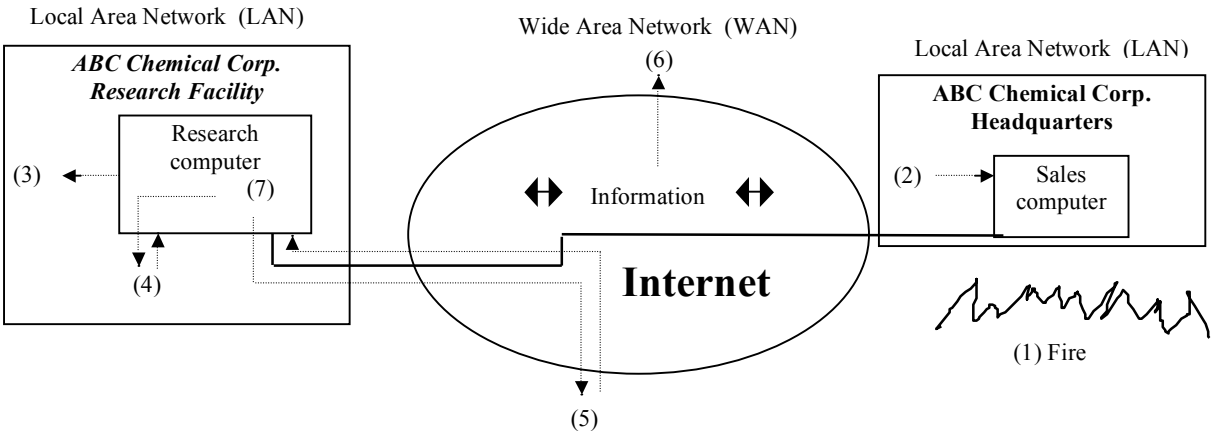

Figure 2. Graphical representation of the seven threat classes

Examples threats in each threat class are given below.

Table 1. Examples of each threat class

Threat class	Example
(1) Non-Human	Fire
(2) Human Error	An employee accidentally deletes a file
(3) Authorized Person	A researcher sells some research data to a competitor
(4) Unauthorized Insider	A mail clerk circumvents the file Access Control mechanisms and obtains some research data
(5) Outsider Attacking Enterprise Data on the LAN	Someone in Iceland breaks into the research computer and obtains research data
(6) Outsider Attacking Enterprise Data Outside the LAN	Someone in Florida intercepts microwave traffic and captures enterprise data that is flowing across the Internet
(7) Malicious Software	Someone in Iowa plants a Trojan horse in a music sharing program

5. HOW THE MODEL SUPPORTS A SCIENTIFIC/SYSTEMATIC TREATMENT OF INFORMATION PROTECTION MATERIAL

The following table lists the threat classes and general categories of protection techniques that address each threat class. Viewed this way it is possible to see how the categories of common Information Protection techniques (Passwords, ACLs, Firewalls, Cryptography, etc.) fit into the big picture of Information Protection. And since these threat classes contain all threats (see section 7 for the proof) complete protection against all threats is achieved if each threat class is adequately addressed.

Table 2. Categories of protection techniques that address each threat class

Threat Class	General Protection Measures
Non-Human (1)	**Physical Security**
Human Error (2)	**Error Security** Training Secure Defaults Simple interfaces
Authorized Person (3)	**Personnel Security** Personnel Security (background checks) Separation of Duties Audit
Unauthorized Insider (4)	**Operating System Security** Identification and Authentication File Access Control Memory Protection (Process Separation, Rings)
Outsider Attacking Enterprise Data on the LAN (5)	**Perimeter Security** Firewalls (Network Traffic Access Control)
Outsider Attacking Enterprise Data Outside the LAN (6)	**Communications Security** Cryptographic Access Control
Malicious Software (7)	**Malicious Software Security** Trojan horses: MAC policies Worms: Assurance techniques Viruses: Ad Hoc Antivirus techniques

Although the assignment of these general categories of protection techniques to the seven threat classes is somewhat consistent it is not perfect. For example, parity bits and backup systems, which are used to address Non-Human threats, are not Physical Security techniques. Similarly, the cryptographic technique of Digital Signatures (which is part of Communications Security) is needed to protect the authenticity of message source information against threat classes 4 and 5. And particularly noteworthy is the omission of Physical Security as a technique for addressing the Unauthorized Insider. If an Insider can get physical access to a server system they can access all its data in spite of Operating System Security (I&A, ACLs and Memory Protection). They simply steal the hard drive or boot the server off of removable media that contains hacking tools.

These exceptions do not detract from the threat model, though. They simply mean that the names of the general categories of protection techniques (Physical Security, Error Security, Personnel Security, etc.) that address the classes of the threat model are not perfect. The names listed above were selected because most of them are common terms (Physical Security, Personnel Security, Perimeter Security, Communications Security) and they are sufficiently correct since they encompass most of techniques that are used to address the threat classes.

Since many Information Protection measures are based on Access Control there are a few supplementary Information Protection techniques that specifically address failures in Access Control mechanisms. For completeness the model needs an additional section that contains these techniques.

5.1 Techniques that address failures in Access Control mechanisms

The table below identifies three techniques that are frequently used to address failures in Access Control mechanisms.

Table 3. Common Techniques for Addressing Access Control Failures

Access Control Faliures	• Design and Development Principles that Promote a Sufficient Degree of Assurance • Defense in Depth • Intrusion Detection

The technique of system Assurance is presented first. When a Web Server program has a Buffer Overflow vulnerability that lets an unauthorized user access the system it is a failure in an Access Control mechanism. This may seem odd because we generally do not consider a Web Server program as part of a system's Access Control mechanisms. It is, however. Implicitly we expect the Web Server program to not grant system access to unauthorized users. This is typically not an explicit policy but it is an implicit policy. When such a program does grant an unauthorized user system access it is failing to enforce this implicit policy and hence it is an instance of a failure in an Access Control mechanism. Thus, design and development techniques that promote assurance or trustworthiness in systems address this type of problem.

Defense in Depth is another technique that can effectively address failures in Access Control Mechanisms. In a Defense in Depth strategy more than one protection mechanism is used in series to protect an asset. Two Access Control mechanisms are in series if access to an asset requires the permission of both mechanisms.

A Defense in Depth strategy is optimal when failures of the individual the Access Control mechanisms are statistically independent events. When this is the case the likelihood of concurrent failures in all Access Control mechanisms is the product of the failure rates of each individual mechanism. Typically this greatly reduces the overall failure rate.

Intrusion Detection is another technique that is commonly used to address the problem of Access Control failures. In general, Intrusion Detection techniques look for activities that are either a result of an Access Control failure or an unauthorized action, such as port probing, that is not easily prevented by Access Control techniques.

6. HOW THE MODEL SUPPORTS A THREAT-FIRST PRESENTATION OF INFORMATION PROTECTION TECHNIQUES

A straightforward application of the threat model to the presentation of Information Protection material could simply allocate a section or chapter to each threat class, such as is done below.

Table 4. Applying the Threat Model to the Presentation of Information Protection Material

Chapters	General Information Protection Measures
1. Non-Human Threat	**Physical Security**
2. Human Error Threat	**Error Security**
3. Authorized Person Threat	**Personnel Security**
4. Unauthorized Insider Threat	**Operating System Security**
5. Outsider Attacking Enterprise Data on the LAN Threat	**Perimeter Security**
6. Outsider Attacking Enterprise Data Outside the LAN Threat	**Communications Security**
7. Malicious Software Threat	**Malicious Software**
8. Access Control Failures	**Access Control Failure Techniques**

In order to instill interest in the student the presentation of each protection technique should begin with the threat that is addressed by the protection technique. Chapter 4 above demonstrates this well.

Hypothesize a system with no user Identification capabilities and show how an Insider user can access information that they are not authorized to access.

Hypothesize a system with a user Identification capability and a file access control mechanism (e.g., ACLs and ACL program) but no user Authentication capability and show how an Insider can access information that they are not authorized to access.

Hypothesize a system with user Identification and Authentication and a file access control mechanism (e.g., ACLs and ACL program) and

- consider storing the password information in the clear,
- consider storing the password information encrypted with conventional cryptography,
- consider storing the password information in a hashed format.
  - Introduce the Dictionary Attack.
  - Introduce the Brute Force Attack.
  - Go over password selection guidelines

Hypothesize a system with user Identification and Authentication (I&A) and a file access control mechanism (e.g., ACLs and ACL program) but no process address space control and show how an Insider can directly read and write bytes of memory and access information that they are not authorized to access.

Hypothesize a system with user Identification and Authentication, a file access control mechanism (e.g., ACLs and ACL program) and process address space control but no ring

mechanism and show how an Insider can bypass the file access controls and ultimately access information that they are not authorized to access.

The conclusion of Chapter 4 is that all aspects of OS Security are necessary to adequately address the Insider threat. As previously mentioned, Physical Security also needs to be mentioned because it too plays a role in addressing the Insider Threat.

This arrangement of material is starkly different from what is commonly done today. Most treatments today have a chapter on I&A, a chapter on file access control policies and mechanisms and a chapter on Operating System Security that presents rings. The arrangement above shows that all these techniques work in concert together to address the same threat, which is the Unauthorized Insider threat.

Chapter 7 (Malicious Software Threat) above also demonstrates the value of this approach. Since Mandatory Access Control (MAC) policies are primarily implemented to address Trojan horse programs in applications it only makes sense to present the threat (Trojan horses) and the protection technique (MAC policies) in the same chapter.

Traditional treatments obscure the connection between MAC policies and Trojan horse programs by presenting MAC policies (along with Discretionary Access Control policies) in a Policies chapter and Trojan horse programs in a Malicious Software chapter. When presented this way the reader has to make the connection between the purpose of a MAC policy and the Trojan horse program threat on their own.

7. WHY THE SEVEN CLASSES INCLUDE ALL THREATS

The follow figures prove that the seven classes of the model do, in fact, include all threats. The proof involves repeatedly partitioning the set of all threats until the 7 classes of the model are derived. First, the set of All Threats is partitioned into **Human-Based threats** and **Non-Human-Based threats**.

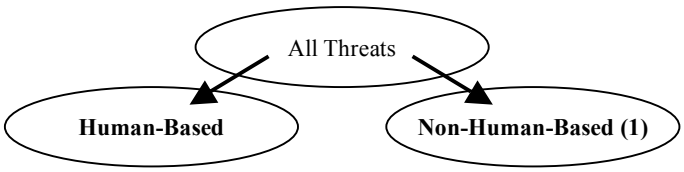

Figure 3. Partitioning all threats into Human and Non-Human-Based threats

Then the set of Human-Based threats is partitioned into the set of **Intentional threats** and **Accidental threats**. Intentional threats are, by definition, performed by Untrustworthy Persons. With the exception of Non-Human threats (fires, lightening, etc.) and Accidental Errors all threats are a result of Untrustworthy Persons.

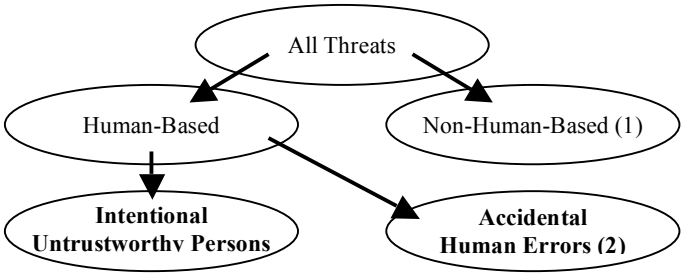

Figure 4. Partitioning Human-Based Threats into Intentional and Accidental Threats

Untrustworthy Persons come in two varieties, those that are authorized to access the data and those that are not authorized to access the data. Thus, the set of Untrustworthy Persons is partitioned into **Unauthorized Persons** and **Authorized Persons**.

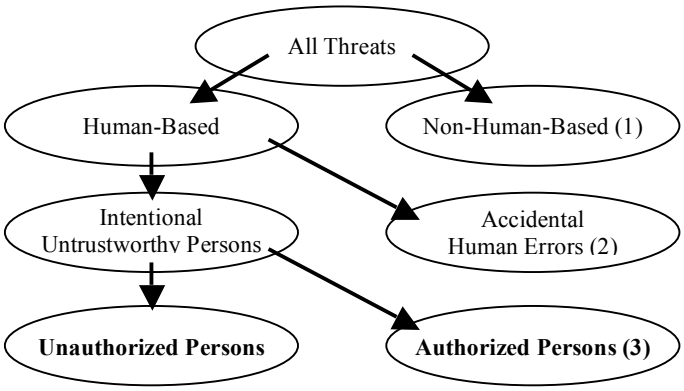

Figure 5. Partitioning Untrustworthy persons into Unauthorized and Authorized sets

Unauthorized Persons come in two varieties, those that are Insiders and those that are Outsiders. Thus, Unauthorized Persons are partitioned accordingly.

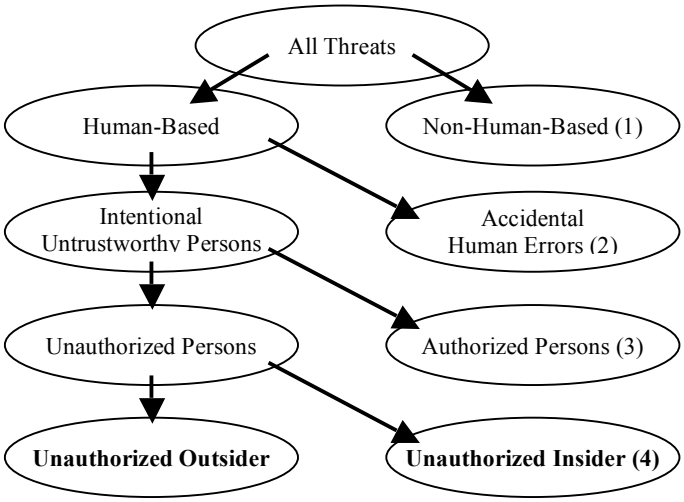

Figure 6. Partitioning Unauthorized Persons into Outsiders and Insiders

Outsiders can threaten enterprise data in two different places, in the environment outside the LAN or on the LAN. Thus, Outsiders are partitioned accordingly.

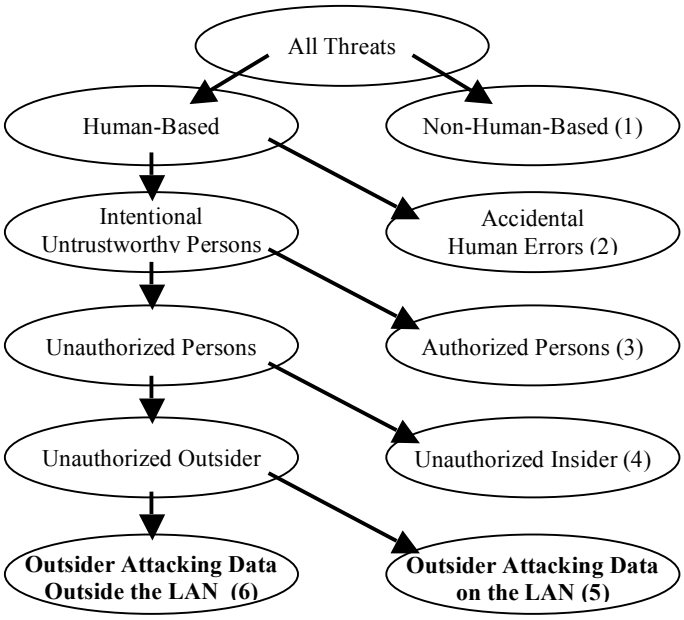

Figure 7. Partitioning Outsiders into LAN Data Attackers and Non-LAN Data Attackers

A seventh threat class (**Malicious Software**) is introduced because Unauthorized Insiders and Outsiders Attacking Enterprise Data on the LAN can attack LAN information in two ways, directly and indirectly. Direct attacks are prevented by appropriate Access Control techniques. This is due to the nature of Access Control, it grants access to authorized persons and denies access to unauthorized persons and both the Unauthorized Insiders and the Outsiders Attacking Data on the LAN are not authorized for the targeted data.

Indirect attacks are generally not addressed by typical Access Control techniques. Indirect attacks use Malicious Software to attack "from the inside." This type of attack is so fundamentally different it is broken out into its own threat class. Access Control techniques keep unauthorized users from gaining direct access, they do not prevent a victim from unknowingly sending their files to a website in Eastern Europe when they execute a Trojan horse program. If the victim is authorized to send data to a website Access Control cannot address this problem.

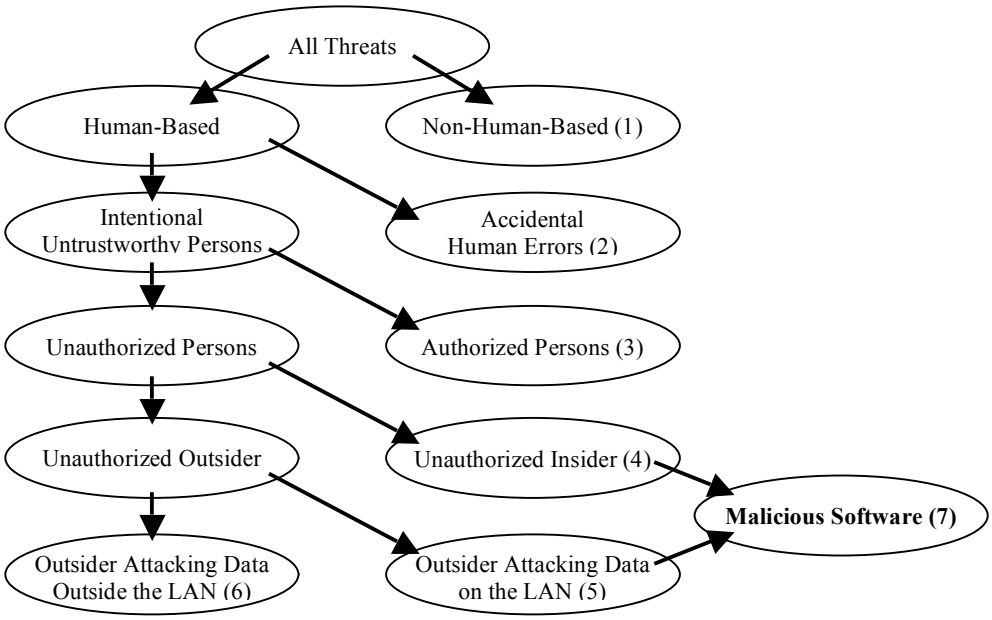

Figure 8. Introducing the Malicious Software Threat Class

Since the seven classes above are derived by partitioning the set of all threats they must contain every information threat. As such, adequately addressing each threat class provides complete coverage against all threats.

8. CONCLUSION

The given threat model provides a systematic way for presenting Information Protection topics. Systematic presentations, like the one proposed here, enhance learning by showing how all the pieces fit together and by establishing a sense of complete coverage against all threats. The use of the threat model presented here also promotes learning because it organizes the material around the threats that provides context and necessity for each protection technique.

ACKNOWLEDGEMENTS

I would like to thank Thuy Nguyen for her valuable comments.

REFERENCES

1. The Merriam-Webster on-line dictionary is at www.m-w.com.